

Serial No. 08/880,616
Art Unit No. 2755

at least one local scheduler associated with each of said more than one computing node comprising means for ascertaining which of said plurality of tasks are assigned tasks, being assigned to each of said plurality of local processes, and means for prioritizing said assigned processes in accordance with said prioritized schedule to allow simultaneous execution of tasks from said more than one application.

REMARKS

The Examiner has again pending Claims 1, 2, 4-16 and 18 under 35 USC § 103 as unpatentable over the teachings of the Cameron patent in view of the admitted prior art (AAPA) regarding the AIX technology. In addition, the Examiner has rejected Claims 3 and 17 under 35 USC § 103 as unpatentable over the teachings of the Cameron patent in view of the AAPA and the teachings of the Ripps reference. Finally, the Examiner has rejected Claims 1 and 11 as unpatentable over Zolnowsky in view of Custer and the AAPA; or, alternatively, as unpatentable over Boland in view of Custer and the AAPA.

Serial No. 08/880,616
Art Unit No. 2755

Applicants wish to first set forth the fundamental differences between the cited prior art and the claimed invention; followed by correlating those differences to the pending claim language. Under the present invention, as taught and claimed, it should be understood that multiple processes per processor are continually waiting for a given task to perform. Traditional OS systems (such as AIX) assign a common priority to all processes and schedule the processes independent of the task importance and independent of other processes in the parallel tasks. In addition, process priority assignment under the prior art does not take into account processes running on other nodes. The present invention, however, provides a method to correlate to execution of tasks -not processes- in time across a parallel machine. The system and taught and claimed has the local scheduler maintain a priority list of ready to execute tasks correlated with local processes. the list is updated by the global scheduler at periodic intervals. The present approach of global (including inter-node) and local scheduling minimizes unused CPU time when an individual task is temporarily blocked or suspended waiting for I/O. Furthermore, the task used to fill idle time under the

Serial No. 08/880,616
Art Unit No. 2755

present invention is the next highest task in importance. Prior art parallel scheduling would either fill the idle cycles with any ready-to-execute process or sit idle until the original task could continue. Therefore, the present invention provides a significant, and non-obvious, advantage over the prior art.

The present invention generalizes scheduling into two levels where Level 1 provides a priority assignment for tasks and level 2 (independently executing on each processing element) provides an association between tasks and processes and schedules the processes optimally to satisfy the global priority arrangement. Ordering between the execution elements will vary between nodes depending upon execution and resource requirements for each particular task to be executed. The result is parallel task ordering for FIFO operations and NOT round-robin scheduling. Actual enforcement is done at level 2 within each processor scheduler. Clearly such is neither taught nor suggested by the Cameron patent which has parallel tasks but single level global scheduling, or by other cited prior art which provides multiple uni-processor tasks on a multi-processor machine with global queuing (Zolnowsky). Applicants assert

Serial No. 08/880,616
Art Unit No. 2755

that the two levels of operation and the two components for executing those two levels of operation are clearly recited in all of the pending claims.

With specific attention to the Cameron patent, the intent and operation of Cameron's scheduler is different from that of the claimed invention. Cameron assumes that each task is evenly divisible across a set of nodes in a partition to utilize such a scheduler. In an embodiment of the present invention relating to database application, the tasks (i.e., queries) are effectively "guaranteed" not to be evenly divisible across a partition. Hence the need for the above-discussed two-level scheduling.

The cited Cameron patent has no means for deciding what process or task should execute when a single process of the currently-scheduled parallel job is suspended or waiting. Since all actions are initiated from the central dispatcher, no provision can be made. Therefore, the CPU is idle for any and all tasks when one is temporarily suspended. Under Cameron, multiple applications can be assigned to a single processor, however, only one can be active and ready to run at a time. (Applicants reiterate and once again direct the Examiner's attention to the statement in Col. 4, line 5 et

Serial No. 08/880,616
Art Unit No. 2755

seq of Cameron that "although more than one application is assigned in a partition, an entire application is scheduled at once across all the nodes on which it is loaded"). Those applications are assigned by the central dispatcher for the partition. The Cameron global scheduler issues a single directive to execute a task at a particular time and, again, has no capability to prioritize tasks or to dynamically assign tasks of multiple processes in order of importance to utilize idle CPU time.

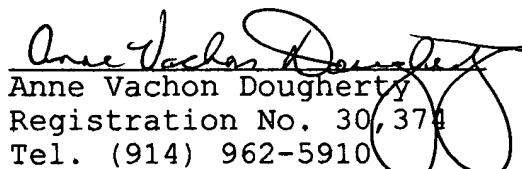
The Examiner had earlier commented that the previously-examined language of independent Claims 1 and 11 did not explicitly recite that the multiple tasks of more than one application could be executed at the same time. Therefore, the language of independent Claims 1 and 11, and therefore also of all of the remaining claims which depend therefrom, was amended by the last-filed response to additionally recite that limitation which was clearly not taught or suggested by the combination of references of Cameron, the AAPA, Ripps, Boland or Zolnowsky. While the Examiner has provided some "response to arguments", Applicants do not understand why the submitted language failed to clarify and distinguish over the cited art.

Serial No. 08/880,616
Art Unit No. 2755

Applicants request that the amended language be specifically addressed in order for a subsequent response to be appropriately formulated. Applicants believe, that, in light of the foregoing explanation, and the earlier-submitted amendments (as well as the current amendments to address the 112 concerns), the currently-pending claims are allowable over the cited art. Withdrawal of the rejections and issuance of the claims it, therefore, respectfully submitted.

Respectfully submitted,

M.A. COHEN, ET AL

By: 
Anne Vachon Dougherty
Registration No. 30,374
Tel. (914) 962-5910